

# Scalable Fine-grained Path Control in Software Defined Networks

Long Luo, Hongfang Yu, Shouxi Luo

Key Laboratory of Optical Fiber Sensing and Communications, Ministry of Education  
University of Electronic Science and Technology of China, Chengdu, P. R. China

**Abstract**—The OpenFlow-based SDN is widely studied to better network performance through planning fine-grained paths. However, being designed to configure path hop-by-hop, it faces the scalability issue—both the flow table overhead and path setup delay are unacceptable for large-scale networks. In this paper, we propose PACO, a framework based on Source Routing to address that problem through quickly pushing paths into the packet header at network edges and pre-installing few rules at the network core. The straightforward implementation of SR is inefficient as it would incur too many path labels; other efficient approaches would sacrifice path flexibility (e.g., DEFO). To implement SR efficiently and flexibly, PACO presents each path as a concatenation of pathlets and introduces algorithms to compute pathlets and concatenate paths with minimum path labels. Our extensive simulations confirm the scalability of PACO as it saves the flow table overhead up to 94% compared with OpenFlow-SDN solutions and show that PACO outperforms SR-SDN solutions by supporting more than 40% paths with few label overhead.

## I. INTRODUCTION AND RELATED WORK

Fine-grained/Flexible routing path control is crucial for today's network to accommodate emerging network services and manage itself [1]–[4]. For security, Qos and traffic engineering purposes, the network operator would perform routine network management tasks including: 1) directing the suspicious traffic (may be from a denial-of-service attack) to a scrubber for inspection, 2) dispersing time-sensitive traffic flow across the low-latency path to meet the time-related Qos requirement, 3) load-balancing bulk traffic flows on under-utilized links to avoid congesting hot/heavy links. Mixing traffic flows over the same path (*i.e.*, shortest path adopted in traditional intra-networks) is problematic as different flows may pitted against each other, thus the network should offer differentiated paths to adapt diverse types of traffic flows.

The OpenFlow-based Software Defined Network (SDN) solution can achieve that goal [5, 6], but it faces a fundamental challenge of scalability [7, 8]. The OpenFlow-SDN approaches achieve flexible path control through configuring hop-by-hop rule for each flow. The scalability issue of such Hop-by-Hop approach reflects in two aspects. One is the insufficient flow table space and the other is the complex path construction process as well the time overhead it incurs. The flow table limitation would prevent network installing routing paths for a large number of traffic flows [9, 10]. The Hop-by-Hop's uncontrollable delay to install a path would prevent the agility of control decision [11, 12]. For example, it would prevent

quick completion of the frequently occurring network updates incurred by network events like failures [12]. Although recent studies have shown that the rule compression technique can reduce the rule overhead to some extent, it is sensitive to the concrete content of rules. Moreover, compressed rules would complicate the update of paths. Therefore, as the network scales up, the Hop-by-Hop solution would find it hard to deal with the increasing demand of routing paths.

To address the above scalability issue, a lot of work resorts to Source Routing (SR) technique to provide network numerous paths [2, 3, 13]–[17]. In the SR-SDN solution, the network pre-installs forwarding rules at switches and encapsulates the entire path in the packet header according to the already configured rules. When there is an ask for a new path, the controller only needs to do the path computation task and control an ingress switch to perform path encapsulation, which is easy even for the large-scale network. Clearly, this approach promises fast path construction as it eliminates the specific rule and operation on hop-by-hop switches. However, there's no such thing as a free lunch—the bits overhead (used to store path) of packet is the price paying for scalability. The straightforward implementation of SR is to directly put labels of each hop along the path into the packet header [13]–[16], which is inefficient as it would incurs too much label overhead, leading to high bandwidth consumption. Another implementation is to have the packet include one or several labels of middlepoints that a path must pass through [3, 17]. This approach adopts shortest paths between two middlepoints, thus it would compromise the path flexibility for lightweight label overhead or suffer from heavy label overhead for flexibility. Others propose to pre-install all the desired paths using state compression algorithm and place a global label for each path in packets [2]. Although the label overhead is lightweight, this approach still meets the scalability issue since it would fail to install all the paths when the condition of its compression algorithm does not be satisfied. Among the state-of-art implementations, we find none is efficient while preserving both the path flexibility and the scalability benefits SR brings.

In this paper, we present PACO, an efficient SR-based approach to provide scalable and flexible path control in software defined networks. Following the principle of SR-SDN solution, PACO restricts the stateful flow/path rule at the network edges and maintains the stateless forwarding rule in the network core. To reduce the overhead of packet, PACO strategically represents

each path as a concatenation of several pathlets/sub-paths<sup>1</sup>. Although the idea of constructing paths with existing pathlets is not new, the key novelty of this paper is proposing an SDN-SR framework together with several accompanied algorithms to implement it in an efficient and flexible way.

To preserve the path flexibility, PACO explicitly programs pathlets and installs them in the network beforehand. Thus, the real challenges PACO tackles are to: 1) decide which pathlets should be pre-installed such that both the path flexibility and scalability are preserved, and 2) how to concatenate a path with minimum pathlets such that the goal of efficiency is achieved.

PACO addresses the first challenge by converting it into a pathlet selection linear programming problem and presenting an efficient pathlet selection algorithm based on our formulated optimization model. As for the path concatenation problem, PACO presents an optimal algorithm based on the branch-and-bound technique and a pathlet nesting scheme.

In summary, we make the following contributions.

**Analysis and Framework.** We disclose the limitations of prior work and show how to use pathlets to overcome them. (section II). We propose a novel SDN-SR based framework to facilitate the scalable and flexible path control. This framework utilizes the global view of SDN control plane to determine the path and follows the forwarding principle of source routing (section III).

**Modeling.** Given the desired paths, we formalize the pathlet selection as a linear programming problem which seeks to compute appropriate pathlets to support all the paths under resource constraints. (section IV).

**Algorithms.** We present a Lagrangian heuristic for the pathlet selection problem and an optimal algorithm for the path concatenation problem. (section IV-F,V).

**Experimental evaluation.** We evaluate our implementation of PACO by simulating extensive path control scenarios for realistic networks. Our results show that PACO hugely outperforms previous techniques in terms of scalability, efficiency and flexibility. (section VI).

## II. MOTIVATION AND RELATED WORK

Fig. 1 shows a scenario where the SDN controller (not depicted for brevity) need to build fine-grained routing paths for the traffic flowing across the controlled network. Provided that there exists four type traffic flows, posing different requirements on their own delivery, between the same ingress-egress pair (see Fig. 1(b)). For the sake of the example, we assume that all the network state (*i.e.*, link bandwidth and delay, traffic type) are acknowledged to the centralized controller (see Fig. 1(a)). Supposing the colored solid arrows represents the programmed optimal flow paths that satisfy the flow constraints (*e.g.*, throughput, latency) and the network constraints (*e.g.*, load balancing, secure data delivery).

In Fig. 1, if the network adopts Hop-by-Hop technique to install the planned paths, it will fail when a rule-scarce node

(say *c*) is capable of adding only 3 more entries to its flow table (while the actual needed is 4). Besides, provided the average time of installing a rule to each hop switch from the controller is  $\delta$ , then the waiting time before a flow enjoying its 5-hops path is  $5\delta$  at worst. Thus, if the lived time of time-sensitive (say flow  $f_1$ ) is less than  $5\delta$ , its data would be dropped at ingress switch before transmission. We say a SDN solution produces a scalable, flexible path control if all the fine-grained paths can be quickly installed in the limited concrete forwarding table.

### A. Previous implementations have limitations

Prior work achieves that goal by exploiting Source Routing (SR) to strictly (*e.g.*, [2, 3, 13]–[17]) or loosely encapsulate the path into each packet header. These SR-SDN approaches hugely mitigates the scalability issue, including flow table overhead and path setup latency, of Hop-by-Hop by moving per-flow routing state from switches to the packet header and pre-installing necessary forwarding rules. According to the encapsulation granularity, we refer current implementations as hop encapsulation (*i.g.*, [13]–[16]), midpoint encapsulation (*i.g.*, [3, 17]) and path encapsulation (*i.g.*, [2]) respectively. Unfortunately, they are inefficient or limited, because they only focus on the benefit but neglect the cost of SR, or sacrifice flexibility and still suffer from scalability.

**Hop encapsulation is inefficient.** It is based on strict SR by inserting labels for the entire path into the packet header at the ingress node ([14]–[16]) or the host side ([13]); each label value indicates the outgoing port number at a hop switch. Since the network diameter is always large in the networks like public WANs or carrier networks, the label overhead of the packet header would be heavy if adopted this technique. This comes with two possible consequences. First, the technique may simply not be applicable if a packet can only carry limited number of labels in the header. For example, a IPv4 packet is able to add at most 4 8-bits labels into its header because of the length constraint of the packet header (while 5 labels is asked to be added for the paths in the Fig. 1). Second, even if multiple labels are supported (*i.e.*, IPv6 extension header), inserting all the port numbers of a path on the packet header would significantly increase the packet size leading to high bandwidth consumption.

**Midpoint encapsulation has limited flexibility.** They are based on loose SR by putting labels of midpoints into each packet header at ingress nodes; each midpoint is a waypoint that the path must go through [3, 17]. The path between two midpoints is computed by the shortest path algorithm of distributed protocols. Thus, this approach may mix requirement-contradictory flows over the same (shortest) path, which may violates the requirements on routing path (*i.e.*,  $f_1$  and  $f_2$  in Fig. 1(b)). To disperse traffic flows over different paths, this approach would introduce more midpoint labels to each packet, bearing the same problems of *hop encapsulation*.

**Path encapsulation cannot always be scalable as expected.** Some proposes to utilize compression approach to pre-install all the desired paths and puts a global label in each packet header to specify which path should be applied [2]. This

<sup>1</sup>In contrast to pathlet routing [18], we use pathlet to represent forwarding paths rather than for routing protocol design.

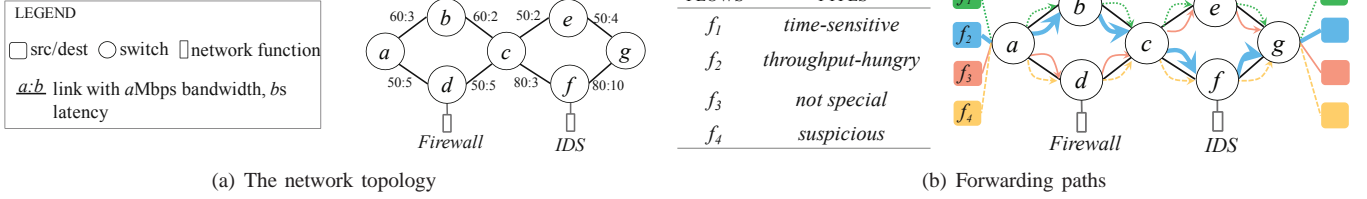


Fig. 1. A forwarding path map of mixed flows.

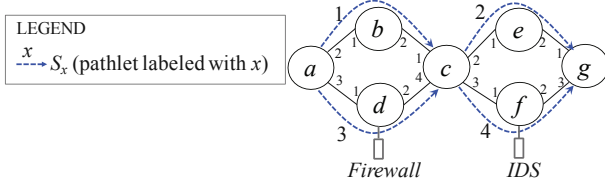


Fig. 2. An example of pathlets.

approach has limited scalability when applying to asymmetrical network topologies since the condition of its state compression algorithm does not be satisfied.

### B. Explicit pathlet is more powerful

The key intuition exploited by PACO is that we can profitably: (1) explicitly define pathlets/sub-paths with the power of SDN, and (2) represent paths as concatenations of pathlets. In the example of Fig. 1, for instance, PACO detects that all the desired paths overlap at pathlets  $(a, b, c)$ ,  $(c, e, f)$ ,  $(a, d, c)$ ,  $(c, g, f)$  (see Fig. 2). Moreover, PACO validates that 1) all the pathlets can be installed in the restricted flowtable, 2) and all the desired paths can be concatenated by pathlets. If applying a label to each pathlet, PACO could represent those paths as concatenations of  $P(f_1) = (1, 2)$ ,  $P(f_2) = (1, 4)$ ,  $P(f_3) = (3, 2)$ ,  $P(f_4) = (3, 4)$  respectively.

**PACO saves much rule space of flow table.** When run on the example in Fig. 1, PACO's rule overhead of installing pathlets in Fig. 2 is much less than that of pushing all the paths to switches (50% rule saving on intermediate switches  $b, d, c, e, f$  compared with Hop-by-Hop approach).

**PACO hugely reduces the number of added labels.** PACO's labels added in each packet of flows (say  $f_i$ ,  $i = 1, \dots, 4$ ) is only 2. This is much more efficient (60% label saving) than *hop encapsulation* that would insert 5 labels into each packet.

**PACO efficiently preserves fine-grained path control.** For the flows of different types, PACO puts their data over granular paths by stitching the explicitly defined pathlets together. Note that the suspicious traffic data (say flow  $f_4$ ) would bypass the firewall if it adopts the *middlepoint encapsulation* approach that would steer its traffic data over the shortest path  $(a, b, c)$  to reach the specified middlepoint (say  $c$  if specified).

## III. PACO OVERVIEW

Fig. 3 overviews PACO. We now describe the formal definition of pathlet PACO adopts, the PACO's framework (section III-A) and the path construction process as well as packet

forwarding process in the context of PACO. We use the terms switch and node interchangeably.

**Pathlet.** A pathlet refers to a directed sub-path  $S = (v_0, v_1, \dots, v_n)$  between node  $v_0$  and  $v_n$  in the controlled network, where  $v_i$  denotes the network node and the nodes are all distinct. It is notice that either the start or the tail node can be the middle node, the ingress node or the egress node. Thus, a pathlet differs much with the end-to-end path in that it is a part of the entire path.

In our PACO model, multiple pathlets are allowed between each node pair for the path flexibility/diversity purpose. Accordingly, there is a need to uniquely identify each pathlet such that the overlapping nodes can perform correct actions instructed by pathlets. To this end, we assign a local *pathlet identifier* or *pid* to each pathlet. Each local identifier is only supported by the nodes over the associated pathlet. The locality of *pid* indicates that the same *pid* can be assigned to multiple disjoint pathlets. Thus, we can use limited bits to express all the *pids*. For convenience of description, we use  $S_i$  to denote the pathlet labeled with identifier  $i$  in the following text, unless explicitly stated otherwise.

### A. Framework

Fig.3 overviews the framework of PACO. As with all the SDN-based approaches, we consider the network includes the controlled forwarding nodes and a logically centralized controller. The controller directly manages the network paths compliant with both the flow and network constraints.

1) *Functionality of a forwarding node:* In our framework, a controlled forwarding node shall be responsible for the following functionality.

**Edge node** provides the fine-grained path for each arriving packet/flow. In particular, the ingress node is responsible for encoding the desired path into the packet header as a list of pathlet identifiers while the egress node would throw packets to corresponding outgoing interfaces.

**Core node** performs simple packet forwarding actions according to the pathlet identifiers encapsulated in packets. Specifically, since the matching field is a pathlet identifier, the look up operation is simple and easy.

2) *Functionality of the controller:* Inside the controller, it has three main functional modules: pathlet manager, path concatenation, and flow table computation. We outline the detail role of each below.

**Pathlet manager** controls all the pathlets in the network, which plays a pivotal role in flexibly provisioning fine-grained paths



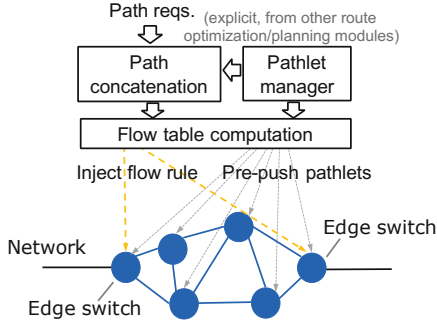


Fig. 3. Framework of PACO.

for all differentiated network services and applications. In the implementation of PACO, *pathlet manager* selects the pathlets that should be installed at the early stage of network configuration. Besides, this module provides the global perspective of pathlets for the *path concatenation* module. (Section IV)

**Path concatenation** provides support for computing the concatenation paths, taking the full end-to-end paths and pathlets as inputs. Indeed, this module focuses only on path concatenation, which is only a small subset of control plane functionality. It is important to clarify that PACO does none of additional route planning related tasks [3, 5, 6]. Considering this module supports the necessary functionality, we leave the route planning to other programs out of consideration for clean separation of network function modularity. Thus, the input only involves the explicit desired paths. For each explicit path request, this module efficiently transforms it into qualified concatenation path through stitching minimum pathlets together. (Section V)

**Flow table computation** synthesizes rules for the pathlet pre-installation and the flow path setup. We here describe the rules at a high-level and do not restrict the concrete implementation.

For a pathlet, the module crafts stateless rules  $r(pid, action)$  for the middle network nodes to match against  $pid$  and apply corresponding  $action$  on matched packet. In particular, a pathlet involves two types of actions:  $forward(v, o)$  and  $pop(v, o)$ . A  $forward(v, o)$  instructs a switch  $v$  to put a matched packet on its outgoing interface  $o$ . While a  $pop(v, o)$  instructs a switch  $v$  firstly pop the top pathlet identifier of a matched packet and then forward that packet to the output  $o$ . In particular, for a pathlet  $S_i = (v_0, v_1, \dots, v_n)$ , PACO computes rules  $r(i, forward(v_k, o))$ , ( $k = 0, \dots, n-2$ ) and a rule  $r(i, pop(v_{n-1}, o))$  for the nodes along that pathlet.

While for a requesting flow, this module crafts a stateful flow rule  $r(fid, action)$  to instruct switches match against the flow identifier  $fid$  (any specified  $k-tuple$ ) and apply action (*insert* or *forward*) on the packets belonging to the matched flow. The action  $insert(v, f, sl, o)$  instructs switch  $v$  firstly insert a list of pathlet identifiers  $sl$  to the packets of flow  $f$  and then forward packets to output  $o$ . Specifically, a flow rule  $r(f, insert(v_{in}, f, sl))$  should be installed at ingress switch  $v_{in}$  and a rule  $r(f, forward(u_{out}, f, o))$  is designed to be placed on egress switch  $u_{out}$ .

**What matters:** When talking about PACO, one would ask

the following two questions: 1) how many and which pathlets should the network pre-install in the network such that all of the desired paths can be concatenated and the flow table capacity is enough, 2) how to concatenate path with pathlets such that the pathlet identifiers or labels added on each packet is always limited. In the below, the first problem is addressed by our optimization model and Lagrangian heuristic (Section IV). The second issue is resolved by our efficient optimal path concatenation algorithm (Section V).

#### IV. PATHLET SELECTION

This section details on how to select pathlets. We first state the pathlet selection problem (section IV-A), and then elaborate on the formulation of a linear programming for it (section IV-B). Finally, to obtain a selection solution within polynomial computational complexity, we show how to decompose problem based on Lagrangian relaxation (section IV-C, IV-D) and present our pathlet selection algorithm based on subgradient algorithm (section IV-E, IV-F).

##### A. The problem description

The *pathlet manager* decides what pathlets should be installed in the dataplane such that all the desired path can be concatenated by them.

**Input.** We assume an set of numerous desired paths  $\mathcal{P}^D$  and a large set of pathlet candidates  $\mathcal{S}^C$ .  $\mathcal{P}^D$  can be computed by the logically centralized controller according to the history log/data, reservation services or optimization goals. Note that, PACO only admits different routing paths between each ingress-egress pair. When determining pathlets, one would try to deduce pathlets according to certain topological characteristics (e.g., overlapping, similarity) among all the desired paths at first thought. However, it is very difficult due to the enormous variations in paths and their topological relationships. To simplify the selection problem, we adopt a large set of pathlet candidates  $\mathcal{S}^C$  and try to select a small group pathlets from that big collection to concatenate all the desired paths.

**Constraints to be meet** including flow table capacity constraint and length of path label constraint.

**C1: Flow table constraint.** The flow table size is a constraint on each network node. A flow table contains entries that specify the forwarding rules for different pathlets, and its capacity is limited. To express this constraint quantitatively, we use  $c_v^{free}$  to denote the free flow table capacity of node  $v \in V$ .

**C2: Length of path label constraint.** The length of label to encode a path is another constraint. The length of path label equals to the number of pathlets in a concatenation path and reflects the label overhead of the packet who applies that concatenation. Such packet overhead may waste precious bandwidth. It is notice that the path label optimization is done in our *path concatenation module*. We here only constrain the maybe pathlet number of a concatenation path to a loose upper bound for the purpose of pre-optimization. To meet this constraint, we adopt  $m^{max}$  to represent the maximum allowed pathlets of concatenation path.

**Output.** The *pathlet manager module* returns the a small group of pathlets  $\mathcal{S}^{select} \subseteq \mathcal{S}^C$  such that the concatenation paths  $\mathcal{P}^{concatenate} \subseteq \mathcal{P}^D$  is maximized. The best solution is to be  $\mathcal{P}^{concatenate} = \mathcal{P}^D$ , which is alway promised by our strategy algorithm as shown in section VI.

### B. The formulation

We model the network as a directed graph  $G = (V, E)$ , where set  $V$  represents the network nodes and set  $E$  models the directed links connecting nodes. From the problem described in Section IV-A, we can easily get the links and nodes that each path or pathlet candidate go through. For each link  $e \in E$ , let  $b_{e,P} = 1$  (resp.  $b_{e,P} = 0$ ) if path  $P$  traverses (resp. does not traverse) through link  $e$  and  $a_{e,S} = 1$  (resp.  $a_{e,S} = 0$ ) if pathlet candidate  $S_i$  passes (resp. does not pass) link  $e$ . Besides, we let  $l_P = \sum_{e \in E} b_{e,P}$  to denote the length of path  $P$  and  $l_S = \sum_{e \in E} a_{e,S}$  to denote that of pathlet candidate  $S$ . Moreover, we let  $d_{v,S}$  denote the rules needed on node  $v$  to install pathlet  $S$ .

We formulate the pathlet selection problem  $\mathbf{P}$  as an *Integer Linear Programming* (ILP). To formulate our problem, we introduce the following decision variables. We define a binary decision variable  $x_{S,P}$  to be 1 if pathlet  $S$  is a component of the concatenation of path  $P$ , and 0 otherwise. for each pathlet candidate in  $\mathcal{S}^C$ , we define the decision variable  $t_S \in \{0, 1\}$  to denote whether it is selected. For each path in  $\mathcal{P}^D$ , we define the binary decision variable  $y_P$  to be 1 if path  $P$  cannot be concatenated and 0 otherwise. To maximize the concatenation paths (minimize the unconcatenated paths), now the problem  $\mathbf{P}$  can be formulated as follows:

$$Z_{IP} = \min \sum_{P \in \mathcal{P}^D} y_P \quad (1)$$

$$\text{Subject to: } \forall e \in E, P \in \mathcal{P}^D : \sum_{S \in \mathcal{S}^C} a_{e,S} \times x_{S,P} \leq b_{e,P} \quad (2)$$

$$\forall P \in \mathcal{P}^D : l_P \times (1 - y_P) = \sum_{S \in \mathcal{S}^C} l_S \times x_{S,P} \quad (3)$$

$$\forall P \in \mathcal{P}^D : \sum_{S \in \mathcal{S}^C} x_{S,P} \leq m^{max} \quad (4)$$

$$\forall S \in \mathcal{S}^C : \sum_{P \in \mathcal{P}^D} x_{S,P} \leq |\mathcal{P}^D| \times t_S \quad (5)$$

$$\forall v \in V : \sum_{S \in \mathcal{S}^C} d_{v,S} \times t_S \leq c_v^{free} \quad (6)$$

$$\forall S \in \mathcal{S}^C, \forall P \in \mathcal{P}^D : x_{S,P} \in \{0, 1\} \quad (7)$$

$$\forall P \in \mathcal{P}^D : y_P \in \{0, 1\} \quad (8)$$

$$\forall S \in \mathcal{S}^C : t_S \in \{0, 1\} \quad (9)$$

Constraints Ineq. (2) and Eq. (3) constrain that the concatenation of a desired path should have the same path structure with it. Constraints Ineq. (4) are length of path label/pathlet constraints on each concatenation path. Ineq. (5) and Ineq. (6) express the flow table constraints on each node.

After solving the above problem, for the desired path  $\mathcal{P}^D$ , we can obtain the maximum number of concatenation paths  $\mathcal{P}^{concatenate}$  (through  $y$ ) and the pathlet selection solution

$\mathcal{S}^{select}$  (through  $t$ ), under the constraints on flow table size and concatenation path length.

To solve the problem  $\mathbf{P}$  within polynomial computational complexity, we design an *pathlet selection heuristic* (Algorithm 1) to obtain a near-optimal solution. Based on decomposition techniques of Lagrangian relaxation and “divide and conquer”, subgradient algorithm, and branch-and-bound technique, the Algorithm 1 can be efficiently implemented. The following shows more details.

### C. Lagrangian Dual Problem

We dualize the constraints in Ineq. (5) to obtain the Lagrangian dual problem as the relaxing problem which can be further decomposed into two sub-problems, each can be solved within polynomial computational complexity.

In particular, we first relax the Ineq. (5) by bringing them to the objective function with associated Lagrangian multiplier vector  $\lambda = \{\lambda_S \geq 0, S \in \mathcal{S}^C\}$ . We get the Lagrangian relaxation problem  $\mathbf{P}_{LR}$  as follows:

$$\begin{aligned} Z_{LR}(\lambda) &= \min \sum_{P \in \mathcal{P}^D} y_P + \sum_{S \in \mathcal{S}^C} \lambda_S \left( \sum_{P \in \mathcal{P}^D} x_{S,P} - |\mathcal{P}^D| t_S \right) \\ &= \min \sum_{P \in \mathcal{P}^D} (y_P + \sum_{S \in \mathcal{S}^C} \lambda_S x_{S,P}) - |\mathcal{P}^D| \sum_{S \in \mathcal{S}^C} \lambda_S t_S \end{aligned} \quad (10)$$

Subject to (2) – (4) and (6) – (9).

For any Lagrangian multiplier  $\lambda$ , the value  $Z_{LR}(\lambda)$  of the Lagrangian function is a lower bound on the optimal objective function value of the original minimization problem  $\mathbf{P}$  [19]. Thus, to obtain the sharpest possible lower bound, we would need to solve the following Lagrangian dual problem  $\mathbf{P}_{LD}$  of  $\mathbf{P}$ :

$$Z_{LD} = \max_{\lambda > 0} Z_{LR}(\lambda) \quad (12)$$

We observe that constraints Ineq. (2)-(4), (7)-(8) only include the group of variables  $\{x_{S,P}, y_P\}$  and Ineq. (6), (9) are only related with the group of variables  $\{t_S\}$ . Therefore, problem  $\mathbf{P}_{LR}$  can be decomposed into two sub-problems involving clean separated variable set,  $\mathbf{P}_{Sub1}$  and  $\mathbf{P}_{Sub2}$  as follows,

$$Z_{Sub1}(\lambda) = \min \sum_{P \in \mathcal{P}^D} (y_P + \sum_{S \in \mathcal{S}^C} \lambda_S x_{S,P}) \quad (13)$$

under constraints Ineq. (2)-(4) and Ineq. (7)-(8), and

$$Z_{Sub2}(\lambda) = \min -|\mathcal{P}^D| \sum_{S \in \mathcal{S}^C} \lambda_S t_S \quad (14)$$

under constraints Ineq. (6) and Ineq. (9).

We observe that problem  $\mathbf{P}_{Sub1}$  can be further decomposed into  $|\mathcal{P}^D|$  mini-problems through “divide and conquer” approach. Finally, we can obtain the optimal solution for  $\mathbf{P}_{Sub1}$  by compositing the solutions of mini-problems. While for the  $\mathbf{P}_{Sub2}$ , it is the general 0-1 linear programming problem, which can be efficiently solved through the exact algorithm of branch-and-bound. In particular, we obtain the optimal solution of the  $\mathbf{P}_{Sub2}$  through the bound-and-bound algorithm in [20].

Given  $\lambda$ , we can compute  $Z_{LR}(\lambda)$  through the optimal objective values of its two sub-problems. In particular, we obtain  $Z_{LR}(\lambda) = Z_{Sub1}(\lambda) + Z_{Sub2}(\lambda)$ .

#### D. Decomposition for sub-problem $\mathbf{P}_{Sub1}$

We adopt the “divide and conquer” approach to further decompose the problem  $\mathbf{P}_{Sub1}$  into  $|\mathcal{P}^D|$  independent mini-problems, each involves a single desired path. For a single path  $P \in \mathcal{P}^D$ , we only need to solve the following optimization problem  $\mathbf{P}_{Sub1,P}$ .

$$Z_{Sub1,j}(\lambda) = \min y_P + \sum_{S \in \mathcal{S}^C} \lambda_S x_{S,P} \quad (15)$$

$$\text{Subject to: } \forall e \in E : \sum_{S \in \mathcal{S}^C} a_{e,S} \times x_{S,P} \leq b_{e,P} \quad (16)$$

$$l_P \times (1 - y_P) = \sum_{S \in \mathcal{S}^C} l_S \times x_{S,P} \quad (17)$$

$$\sum_{S \in \mathcal{S}^C} x_{S,P} \leq m^{\max} \quad (18)$$

$$y_P \in \{0, 1\}, \forall S \in \mathcal{S}^C : x_{S,P} \in \{0, 1\} \quad (19)$$

This above problem can be efficiently solved by adopting the branch-and-bound technique.

After obtaining the optimal objective value  $Z_{Sub1,P}(\lambda)$  of problem  $\mathbf{P}_{Sub1,P}$ ,  $P \in \mathcal{P}^D$ , we can easily compute  $Z_{Sub1}^\lambda = \sum_{P \in \mathcal{P}^D} Z_{Sub1,P}(\lambda)$ .

#### E. Multiplier selection and update

It is clear that the selection of multipliers  $\lambda$  is important for obtaining a good quality or tightness lower bound of problem  $\mathbf{P}$ . We start with an initial multiplier vector and adopt the subgradient optimization to iteratively update  $\lambda$  because of its well-known efficiency.

We choose a starting point  $\lambda^0 = 10^{-3}$ . While at the  $(k+1)th$  iteration, multiplier vector  $\lambda^{k+1}$  can be obtained by  $\lambda^{k+1} = \max\{\lambda^k + \theta_k \mu^k, 0\}$ , in which  $\lambda^k$ ,  $\theta_k$ ,  $\mu^k$  respectively denotes the multiplier vector, the step size, the subgradient vector used in the  $kth$  iteration. Their detail calculations are as follows:

- $\mu^k$ : The subgradient vector  $\mu^k$  is computed by  $\mu_S^k = \sum_{P \in \mathcal{P}^D} x_{S,P}^k - |\mathcal{P}^D| t_S^k$ ,  $S \in \mathcal{S}^C$ , where  $x_{S,P}^k$  and  $t_S^k$  respectively denotes the values of variables  $x_{S,P}$  and  $t_S$  in the optimal solution of problem  $\mathbf{P}^\lambda$  obtained at  $\lambda^k$ .
- $\theta_k$ : The positive step size  $\theta_k$  can be acquired by a common method in [19] as follows:  $\theta_k = \frac{\beta_k(z_{UP}^k - z_{LB}^k)}{\|\mu^k\|^2}$ , where  $\beta_k$  is a positive scalar satisfying  $0 < \beta_k \leq 2$  and  $z_{UP}^k$  (resp.  $z_{LB}^k$ ) denotes an upper (resp. lower) bound on the optimal objective of the original optimization programming at iteration  $k$ . To obtain a lower bound closer to the optimal value, we compute  $z_{LB}^k$  by  $\max\{B^{\lambda^k}, z_{LB}^{k-1}\}$ , where  $B^{\lambda^k}$  is the objective value of the Lagrangian relaxation programming at  $\lambda^k$ . Besides, the objective value of a feasible solution of the original problem is obviously an upper bound on the optimal objective of the primal problem. Thus, to obtain the smaller upper bound, the  $z_{UP}^k$  is computed by  $\min\{z_{FE}^k, z_{UP}^{k-1}\}$ , where  $z_{FE}^k$  respectively denotes the objective value of a feasible solution at iteration  $k$ .

- $z_{FE}^k$ : To construct a primal feasible solution  $z_{FE}^k$ , we let  $t^k = \{t_S^k, S \in \mathcal{S}^C\}$  be known parameters and solve the optimization programming with only constraints Ineq. (2)-Ineq. (5). It is clearly that the value of  $t^k$  have already satisfied constraints Ineq. (6). Therefore, the obtained solution obviously satisfies the feasibility.

#### F. Pathlet selection heuristic (Algorithm 1) based on Lagrangian relaxation

We now describe our Algorithm 1 which is designed based on Lagrangian relaxation and subgradient optimization to solve the original programming. Specifically, sub-problems  $\mathbf{P}_{Sub1}$  and  $\mathbf{P}_{Sub2}$  are solved with multiplier vector  $\lambda^k$  at iteration  $k$ . At each iteration  $k$ , we can obtain a feasible solution for the original problem based on  $t^k$  and an upper bound of the original programming. We maintain an upper bound  $z_{UP}^k$  as the smallest upper bound we have obtained within  $k$  iterations. On the other hand,  $z_{LB}^k$  denotes the maximum value of the objective of problem  $\mathbf{P}_{LR}$  after  $k$  iterations, which is a lower bound the the original programming.

To obtain a satisfactory solution within limited computation time, the Algorithm 1 is stopped when one of the conditions is satisfied: 1) the number of iteration  $k$  reaches the iteration limit  $|\mathcal{P}|$ ; 2) the difference between  $z_{LB}^k$  and  $z_{UP}^k$  is less than a threshold  $\epsilon^*$ ; 3) the lower bound does not increase for more than a number of iterations  $T'$ . After the algorithm is terminated, it returns the feasible solution  $\pi^*$ , which reaches to the minimum upper bound during the iterations.

In very rare cases, some of the desired paths cannot be concatenated after running the Algorithm 1 only once. This is because of the improper input of pathlet candidates  $\mathcal{S}^C$ . To maximize the concatenation paths, we present algorithm 2 that runs Algorithm 1 repeatedly until all the paths can be concatenated (see Line 5-17 in Algorithm 2). Specifically, at each running time, Algorithm 1 takes only the unconcatenated paths and a set of random pathlet candidates as input.

### V. PATH CONCATENATION

As described in section III-A, the task of *path concatenation module* is to concatenate each on-demand path with the pathlets selected by *segment manager module*.

#### A. Problem description

As we have described previous, as the path is encoded in each packet, the path label overhead should be minimized for saving bandwidth resource. Therefore, for a on-demand path  $P$ , to minimize the path label overhead, the problem here is to compute the minimum pathlets from the selected pathlets  $\mathcal{S}^{select}$  to construct the concatenation for it. We say a concatenation solution is optimal for a path when it has the minimum pathlets. Although the pathlet selection algorithm in section IV has bounded the number of pathlets of a concatenation path, it is a loose constraint and does not promises the optimal concatenation for each single path.

---

**Algorithm 1** OneRoundSelection( $\mathcal{S}^C, \mathcal{P}^D$ )

---

```
1:  $k \leftarrow 1$  and  $t' \leftarrow 0$ ;  
2:  $z_{UP}^0 \leftarrow +\infty$  and  $z_{LB}^0 \leftarrow -\infty$ ;  $\beta_1 \leftarrow 2$  and  $\epsilon^1 \leftarrow +\infty$   
3: Let  $\lambda_S^1 \leftarrow 10^{-3}$ ,  $S \in \mathcal{S}^C$ ;  
4: while  $k < |\mathcal{P}^D|$  and  $\epsilon^t > \epsilon^*$  and  $t' < T'$  do  
5:   Solve problem  $\mathbf{P}_{Sub1}$ ; Obtain  $Z_{Sub1}^k$  and  $x_{S,P}^k, S \in \mathcal{S}^C, P \in \mathcal{P}^D$ ;  
6:   Solve problem  $\mathbf{P}_{Sub2}$ ; Obtain  $Z_{Sub2}^k$  and  $t_S^k, S \in \mathcal{S}^C$ ;  
7:    $Z_{LR}^k = Z_{Sub1}^k + Z_{Sub2}^k$ ;  
8:   Let  $t^k$  be an given parameters and solve the problem  $\mathbf{P}$  with equations (2)-(5), let  $z_{FE}^k$  be the objective value of the solving problem;  
9:   Let  $z_{UP}^k = \min\{z_{FE}^k, z_{UP}^{k-1}\}$ ;  
10:  if  $z_{UP}^k < z_{UP}^{k-1}$  then  
11:    Record the feasible solution  $R^*$ ;  
12:  end if  
13:  Let  $z_{LB}^k = \max\{B^k, z_{LB}^{k-1}\}$ ;  
14:  if  $z_{LB}^k > z_{LB}^{k-1}$  then  
15:     $t' = 0$ ;  
16:  else  
17:     $t' = t' + 1$ ;  
18:  end if  
19:  if  $t' \geq 4$  then  
20:     $\beta_k = \beta_{k-1}/2$ ;  
21:  end if  
22:   $\epsilon^{k+1} = z_{UP}^k - z_{LB}^k$ ;  
23:  // update the multipliers  $\lambda$   
24:   $\mu_S^k = \sum_{P \in \mathcal{P}^D} x_{S,P}^k - |\mathcal{P}^D| t_S^k, S \in \mathcal{S}^C$ ;  
25:   $\theta^k = \frac{\beta_k(z_{UP}^k - z_{LB}^k)}{\|\mu^k\|^2}$ ;  
26:   $\lambda_S^{k+1} = \max\{\lambda_S^k + \theta^k \mu_S^k, 0\}, S \in \mathcal{S}^C$ ;  
27:   $k = k + 1$ ;  
28: end while  
29: return  $R^*$  and  $z_{UP}^{k-1}$ ;
```

---

---

**Algorithm 2** SelectPathlet( $\mathcal{P}^D$ )

---

```
1: //  $\mathcal{P}^D$ : all the desired paths  
2:  $\mathcal{S}^{select} \leftarrow \emptyset$ ; ▷ The selected pathlets  
3:  $\mathcal{P}^{concatenate} \leftarrow \emptyset$ ; ▷ The concatenation paths  
4:  $n \leftarrow 1$ ;  
5: while  $n \leq N$  do ▷  $N$  is the iteration limit  
6:    $\mathcal{P}^{non} = \mathcal{P}^D - \mathcal{P}^{concatenate}$ ;  
7:   Generate pathlet candidates  $\mathcal{S}^C$ ; ▷ Randomly  
8:    $R^* = \text{SELECTION}(\mathcal{S}^C, \mathcal{P}^{non})$ ;  
9:   Extract the concatenation paths  $\mathcal{P}_n$  and the selected pathlets  $\mathcal{S}_n$  from  $R^*$ ;  
10:   $\mathcal{P}^{concatenate}.\text{addPath}(\mathcal{P}_n)$ ;  
11:   $\mathcal{S}^{select}.\text{addPathlet}(\mathcal{S}_n)$ ;  
12:  if  $\mathcal{P}^{concatenate} = \mathcal{P}^D$  then  
13:    return  $\mathcal{S}^{select}$ ;  
14:  end if  
15:  Update  $c_v^{free}, v \in V$  according to  $\mathcal{S}^{select}$ ;  
16:   $n \leftarrow n + 1$ ;  
17: end while  
18: return  $\mathcal{S}^{select}$ ;
```

---

**B. Problem analysis**

Since there is no competition for network resources when actually concatenating the requesting paths, we focus on the problem of processing one on-demand path. If there are multiple path requests, we can compute their concatenation path in parallel. In particular, we exploit the fact that computing the optimal solution for a single path is not NP-hard:

**Theorem 1** An optimal concatenation for a given request  $P$  can be computed in polynomial time.

*Proof.* The proof is constructive. We first prune all pathlets  $S \in \mathcal{S}^{select}$  whose traversing links do not be traversed by  $P$ . Adopting the above pruning method, the precise pathlet candidates  $\mathcal{S}^P$  for  $P$  is obtained, which has been proved to be  $|\mathcal{S}^P| \leq (\frac{l_P^2 + l_P - 2}{2})$ . For each  $S \in \mathcal{S}^P$ , since  $P$  goes through all links  $e$  on it, thus it is a proper subset of  $P$ . Moreover, the number of pathlets used to concatenate  $P$  must be less than  $l_P$ . Therefore, we can compute all the possible combinations  $\mathcal{P}^{enumerate}$  ( $|\mathcal{P}^{enumerate}| = \sum_{i=2}^{l_P-1} C_{\mathcal{S}^P}^i$ ) of pathlets in  $\mathcal{S}^P$ . Since the number of possible combinations is limited, we can obtain the optimal one in polynomial time even adopting the enumeration approach, proving the theorem.

Theorem 2 is an important instruction for our path concatenation algorithm design, as it allows us to devise algorithm that always returns optimal concatenation. However, the natural enumeration approach that verifies all the possible combinations is not efficient.

**C. Algorithm**

We present an efficient path concatenation algorithm (Algorithm 3) that strategically combines incremental enumeration with “First-Fit” to obtain the optimal concatenation path (see Algorithm 3). Firstly, given the selected pathlets  $\mathcal{S}^{select}$  and a requesting path  $P$ , Algorithm 3 extracts the precise pathlet candidates for  $P$ . Then, Algorithm 3 performs the following three steps to compute the optimal concatenation path for  $P$ . Step 1: *Enumerate* all the combinations with  $m$  pathlets, where  $m$  is initialized with 2 (Line 5). Step 2: *Verify* the enumerated combinations obtained from step 1 one by one until finding a satisfactory one (Line 6-10). If this step can not find a satisfactory concatenation path under the current  $m$ , it will then increase the  $m$  by 1 and go to step 1. Algorithm 3 repeats the step 1, 2 until obtaining a feasible concatenation path  $P_{optimal}^{con}$  or reaching the possible maximum value of  $m$  ( $l_P$ ). Step 3: *Nesting* the obtained  $P_{optimal}^{con}$  if the pathlets included in it exceed the allowed maximum value  $m^{max}$  (Section V-D). Otherwise, Algorithm 3 returns the optimal solution  $P_{optimal}^{con}$ .

**D. Pathlet nesting scheme**

In rare cases, a concatenation would still include many pathlets. This is because the tasks of pathlet selection and concatenation computation are done separately. To handle such case, we adopt a nesting scheme to let each packet carry limited pathlet identifiers at each hop switch. The main idea is simple. For long concatenations, we design representative pathlet to denote several pathlets. The detail of a representative pathlet



---

**Algorithm 3** ConstructPath( $\mathcal{S}^{select}, P$ )

---

```

1: Extract the pathlet candidates  $\mathcal{S}^P$  from selected pathlets
    $\mathcal{S}^{select}$  for desired path  $P$ ;
2:  $P_{optimal}^{con} \leftarrow \emptyset$ ;  $\triangleright$  Store the optimal concatenation path.
3:  $m \leftarrow 2$ ;
4: while  $m \leq l_P$  and  $P_{optimal}^{con} = \emptyset$  do
5:    $\mathcal{P}^{enumerate} \leftarrow Enumeration(\mathcal{S}^P, m)$ ;
6:   for  $P' \in \mathcal{P}^{enumerate}$  do
7:     if  $P' = P$  then
8:        $P_{optimal}^{con} = P'$ , Break;
9:     end if
10:  end for
11:   $m \leftarrow m + 1$ ;
12: end while
13: if  $m \leq m^{max}$  then
14:   return  $P_{optimal}^{con}$ ;
15: else
16:    $P_{optimal}^{con} \leftarrow Nesting(P_{optimal}^{con})$ , and return  $P_{optimal}^{con}$ ;
17: end if

```

---

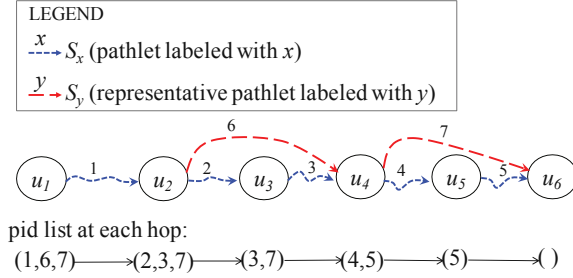


Fig. 4. An example of shortening a long concatenation path.

is invisible at irrelevant nodes and only to be unfolded at the start point of that representative pathlet.

Consider the concatenation of a path  $P_{u_1, u_6}$  is  $\{S_1, S_2, S_3, S_4, S_5\}$  in Fig. 4. As stated in Section III, each packet employing  $P_{u_1, u_6}$  should carry five labels  $sl = \{1, 2, 3, 4, 5\}$  at ingress switch  $u_1$ . The labels can be reduced to three if adopting two representative pathlets:  $S_6$  (for pathlets  $S_2$  and  $S_3$ ) and  $S_7$  (for pathlets  $S_4$  and  $S_5$ ). More generally, PACO can always compute a concatenation with very limited pathlets and pathlet representatives for each path.

## VI. PERFORMANCE EVALUATION

We evaluate PACO by performing 5,000 simulation experiments. In each experiment, we generate a large-scale of desired paths and run our PACO on it. Besides, we collect the rules used to install the pathlets selected by PACO, the paths that can be concatenated by the selected pathlets of PACO and the length of path label of each path.

TABLE I. Experimental Topologies

Topology name	rf3967	rf1755	rf1221	rf6164	rf3257
Nodes	79	87	104	138	161
Links	294	322	302	744	656

As dataset, we adopt the inferred topologies from the Rocketfuel project [21]. To generate large-scale desired paths, we

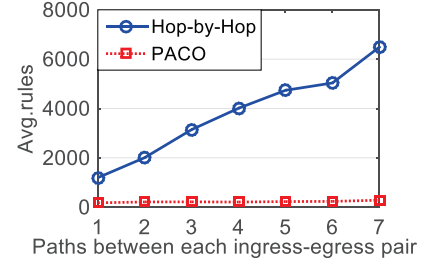


Fig. 5. The rules needed by PACO do not increase with the number of paths.

assume four type of flows (protected, suspicious, bulk and time-sensitive) and compute policy path for each type of flows. In particular, we compute two disjoint paths for a protected flow, a path passing through a random waypoint for a suspicious flow, a bandwidth-sufficient path for a bulk flow and a low-latency path for a time-sensitive flow. In each experiment, we generate  $m$  flows, each of random type, between each node pair and take the flow paths between all node pairs as the desired paths. To be fair, we only admit different paths between the same node pair. As for the pathlet candidates, we generate  $k$ -simple paths between every node pair to make that.

TABLE II. Rule overhead of PACO and that of Hop-by-Hop technique

ISPs	Paths	Max.rules		Avg.rules		$R_{Ave\text{save}}$
		Hop-by-Hop	PACO	Hop-by-Hop	PACO	
rf3967	26,636	9,281	467	2,538	343	86.48%
rf1755	32,373	15,928	613	2,982	268	90.89%
rf1221	41,132	21,970	641	3,095	209	93.25%
rf6164	87,594	31,677	742	4,516	370	91.81%
rf3257	109,516	45,316	926	4,872	286	94.13%

**PACO is scalable.** PACO hugely reduces the number of rules compared with Hop-by-Hop/OpenFlow-SDN technique. Table II shows the results of experiments over a scenario there are four flows between each node pair. The results show PACO can save up to about 94% average rules that would be used by Hop-by-Hop technique. More interesting, the maximum rule number on a node is always less than 1000, which indicates all of the pathlets selected by PACO could be installed in the commodity switch<sup>2</sup>. Besides, we rerun PACO over scenarios of increasing paths. The results of experiments on topology rf1221 depicted in Fig. 5 show the average number of rules of PACO is always less than 300 but that of Hop-by-Hop increases linearly with the paths. This result indicates that *PACO is scalable as the required rules would not increase with the number of paths.*

**PACO always concatenates all of the desired paths flexibly** with the pathlets selected by our algorithm in each and every experiment. We fix the maximum number of pathlets used to concatenate a path to three, as [3] states it can get significant TE improvement with that setting. We respectively run DEFO and PACO over the same set of desired paths. Results are displayed in Fig. 6. *PACO's 100% concatenation indicates the preserving of path flexibility and marks an important difference with previous techniques based on middlepoint encapsulation.*

<sup>2</sup>The commodity switches have at least 2K capacity.



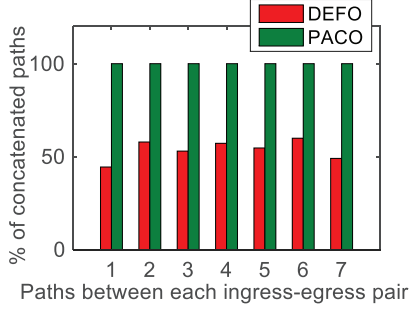


Fig. 6. The concatenated paths of PACO and of DEFO [3].

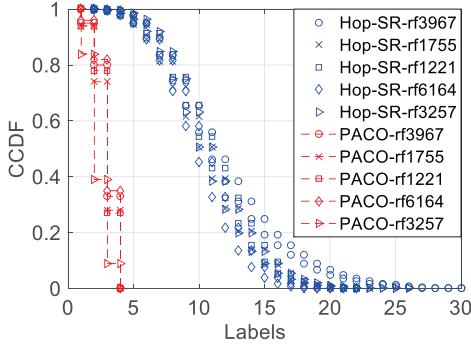


Fig. 7. The length of path label.

The results show that DEFO cannot concatenate more than  $\approx 60\%$  desired paths of the experiments on any topology. In contrast, PACO computes concatenations for all the desired paths in our experiments. This is because PACO adopts much flexible pathlets between two middlepoints rather than shortest paths. The results of DEFO indicate that a set of shortest paths might fit well for a specific goal (*i.e.*, TE) but would struggle to reach other goals (*i.e.*, service function enforcement).

**PACO only needs to adopt limited labels** to encode a path. We evaluate the result of path labels for each experiment of each topology. Fig. 7 depicts the Complementary CDF of the results. A data point  $(x, y)$  in the figure indicates that for a fraction of  $y$  of the desired paths need at least  $x$  path labels. The results highlight that the PACO can concatenate all the paths with at most 4 path labels ( $x = 5, y = 0$ ) across experiments on every topology. While the Hop-SR/encapsulation approach needs at least 10 path labels ( $x = 10$ ) to encode about 45% paths ( $y = 0.45$ ). Indeed, for the Hop-SR method, the number of path labels increases linearly with the path length. *The results indicate that PACO is efficient that the label overhead of all the paths is always lightweight*, while that of Hop-SR/encapsulation is heavy.

## VII. CONCLUSION

In this paper, we present PACO, an Source Routing (SR) based framework to provide scalable and fine-grained path control for SDN networks. In the design, PACO would pre-install pathlets in the network and represents each on-demand path as a concatenation of pathlets. Each packet adopts the forwarding technique of SR. To reduce the label overhead SR brings while preserving path flexibility, we first design

and implement a Lagrangian heuristic to determine which pathlets should be installed in the small flow table such that the all the fine-grained paths can be concatenated. And then we design optimal algorithm to computes the minimum pathlets to represent each path efficiently. The results of evaluation show that PACO outperforms previous approaches: In our experiments, it saves more than 94% rules needed by Hop-by-Hop technique with very limited label overhead and supports 40% more fine-grained paths than DEFO.

## REFERENCES

- [1] S. Vissicchio, O. Tilmans, L. Vanbever, and J. Rexford, "Central control over distributed routing," in *SIGCOMM*, 2015.
- [2] S. Hu and K. e. a. Chen, "Explicit path control in commodity data centers: Design and applications," in *NSDI*, 2015.
- [3] R. Hartert, *et al.*, "A declarative and expressive approach to control forwarding paths in carrier-grade networks," in *SIGCOMM*, 2015.
- [4] S. K. Singh, T. Das, and A. Jukan, "A survey on internet multipath routing and provisioning," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, pp. 2157–2175, 2015.
- [5] S. Jain, A. Kumar, Mandal *et al.*, "B4: Experience with a globally-deployed software defined wan," in *SIGCOMM*, 2013.
- [6] C.-Y. Hong and S. e. a. Kandula, "Achieving high utilization with software-driven wan," in *SIGCOMM*, 2013.
- [7] J. McCauley *et al.*, "Recursive sdn for carrier networks," <https://arxiv.org/abs/1605.07734>.
- [8] S. Yeganeh, A. Tootoonchian, and Y. Ganjali, "On scalability of software-defined networking," *IEEE Communications Magazine*, vol. 2, no. 51, pp. 136–141, 2013.
- [9] S. Luo, H. Yu, and L. Li, "Practical flow table aggregation in sdn," *Computer Networks*, vol. 92, pp. 72–88, 2015.
- [10] N. Katta, O. Alipourfard, J. Rexford *et al.*, "Cacheflow: Dependency-aware rule-caching for software-defined networks," in *SOSR*, 2016.
- [11] K. He *et al.*, "Measuring control plane latency in sdn-enabled switches," in *SOSR*, 2015.
- [12] L. Jin *et al.*, "Dynamic scheduling of network updates," in *SIGCOMM*, 2014.
- [13] X. Jin, N. Farrington, and J. Rexford, "Your data center switch is trying too hard," in *SOSR*, 2016.
- [14] S. A. Jyothi, M. Dong, and P. Godfrey, "Towards a flexible data center fabric with source routing," in *SOSR*, 2015.
- [15] A. e. a. Abujoda, "Sdn-based source routing for scalable service chaining in datacenters," in *IFIP WWIC*, 2016.
- [16] M. Soliman, B. Nandy, I. Lambadaris, and P. Ashwood-Smith, "Exploring source routed forwarding in sdn-based wans," in *ICC*, 2014.
- [17] C. Filsfilis, N. K. Nainar, C. Pignataro, J. C. Cardona, and P. Francois, "The segment routing architecture," in *GLOBECOM*, 2015.
- [18] P. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," in *HOTNETS*, 2008.
- [19] M. L. Fisher, "The lagrangian relaxation method for solving integer programming problems," *Management science*, vol. 50, no. 12\_supplement, pp. 1861–1871, 2004.
- [20] S. Martello and P. Toth, *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.
- [21] N. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," in *SIGCOMM*, 2002.